Red Hat Enterprise Linux OpenStack Platform 5
Release Notes for RHEL7.1 and RHEL6.6 Release

Release details for Red Hat Enterprise Linux OpenStack Platform 5

Red Hat Enterprise Linux OpenStack Platform Documentation Team

Red Hat Enterprise Linux OpenStack Platform 5 Release Notes for RHEL7.1 and RHEL6.6 Release

Release details for Red Hat Enterprise Linux OpenStack Platform 5

Red Hat Enterprise Linux OpenStack Platform Documentation Team
Red Hat Customer Content Services
rhos-docs@redhat.com

Legal Notice

Copyright © 2015 Red Hat, Inc.

This document is licensed by Red Hat under the Creative Commons Attribution-ShareAlike 3.0 Unported License. If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

The Release Notes document the major features, enhancements, and known issues of the Red Hat Enterprise Linux OpenStack Platform 5 release on Red Hat Enterprise Linux 6.6 and 7.1.

Table of Contents

Chapter 1. Product Introduction **2**
1.1. About this Release 2
1.2. Operating System Requirements 3
1.3. RHN/CDN Channels 3
1.4. Product Support 7

Chapter 2. Release Information **9**
2.1. Enhancements 9
2.2. Technology Preview 32
2.3. Release Notes 33
2.4. Known Issues 34
2.5. Deprecated Functionality 37

Chapter 3. Upgrading **39**

Revision History **40**

Chapter 1. Product Introduction

Red Hat Enterprise Linux OpenStack Platform provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

The current Red Hat system is based on OpenStack Juno, and packaged so that available physical hardware can be turned into a private, public, or hybrid cloud platform including:

- Fully distributed object storage
- Persistent block-level storage
- Virtual-machine provisioning engine and image storage
- Authentication and authorization mechanism
- Integrated networking
- Web browser-based GUI for both users and administration.

The Red Hat Enterprise Linux OpenStack Platform IaaS cloud is implemented by a collection of interacting services that control its computing, storage, and networking resources. The cloud is managed using a web-based interface which allows administrators to control, provision, and automate OpenStack resources. Additionally, the OpenStack infrastructure is facilitated through an extensive API, which is also available to end users of the cloud.

Note

For an overview of the OpenStack components and their interfaces, see the "Component Overview"(https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/).

1.1. About this Release

This release of Red Hat Enterprise Linux OpenStack Platform is based on the OpenStack "Icehouse" release. It includes additional features, known issues, and resolved issues specific to Red Hat Enterprise Linux OpenStack Platform.

Only changes specific to Red Hat Enterprise Linux OpenStack Platform are included in this release notes document. The release notes for the OpenStack "Icehouse" release itself are available at the following location:

OpenStack "Icehouse" Release Notes

https://wiki.openstack.org/wiki/ReleaseNotes/Icehouse

Red Hat Enterprise Linux OpenStack Platform uses components from other Red Hat products. Specific information pertaining to the support of these components is available at:

https://access.redhat.com/site/support/policy/updates/openstack/platform/

To evaluate Red Hat Enterprise Linux OpenStack Platform, sign up at:

http://www.redhat.com/openstack/.

Note

The Red Hat Enterprise Linux High Availability Add-On is available for Red Hat Enterprise Linux OpenStack Platform use cases. See the following URL for more details on the add-on: http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/. See the following URL for details on the package versions to use in combination with Red Hat Enterprise Linux OpenStack Platform: https://access.redhat.com/site/solutions/509783

1.2. Operating System Requirements

This version of Red Hat Enterprise Linux OpenStack Platform 5 is supported on the following platforms:

- Red Hat Enterprise Linux 6.6
- Red Hat Enterprise Linux 7.1

To view the entire documentation suite for Red Hat Enterprise Linux OpenStack Platform, see:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform

1.2.1. Web Browser Requirements

The Red Hat Enterprise Linux OpenStack Platform dashboard is a web-based interface that allows you to manage OpenStack resources and services. The dashboard for this release supports the latest stable versions of the following web browsers:

- Chrome
- Firefox
- Firefox ESR
- Internet Explorer 11 and later (with *Compatibility Mode* disabled)

1.3. RHN/CDN Channels

This section discusses channel and repository settings required for deploying Red Hat Enterprise Linux OpenStack Platform 5.

The Red Hat Enterprise Linux OpenStack Platform Installer deploys Red Hat Enteprise Linux OpenStack Platform 5 on Red Hat Enterprise Linux 7.1. However, Red Hat Enterprise Linux OpenStack Platform Installer is only supported on Red Hat Enterprise Linux 6.6. As such, to use this Installer you will have to install it first on a Red Hat Enterprise Linux 6.6 system. From that system, you can then use Red Hat Enterprise Linux OpenStack Platform Installer to deploy Red Hat Enterprise Linux OpenStack 5.0 on a separate Red Hat Enterprise Linux 7.1 installation within the network.

Warning

Although older Red Hat OpenStack repositories are available, you must ensure that your system can no longer access them before installing Red Hat Enterprise Linux OpenStack Platform 5. For example, for CDN, unsubscribe from or disable the following:

- Red Hat OpenStack 1.0 (Essex) -- rhel-server-ost-6-preview-rpms
- Red Hat OpenStack 2.1 (Folsom) -- rhel-server-ost-6-folsom-rpms
- Red Hat Enterprise Linux OpenStack Platform 3 (Grizzly) -- rhel-server-ost-6-3-rpms
- Red Hat Enterprise Linux OpenStack Platform 4 Beta (Havana) -- rhel-6-server-openstack-beta-rpms
- Red Hat Enterprise Linux OpenStack Platform 4 (Havana) -- rhel-6-server-openstack-4.0-rpms

Note

The Red Hat Common for RHEL Server channel is recommended for use if creating custom Red Hat Enterprise Linux guest images that require cloud-init.

For Red Hat Enterprise Linux 6, run:

```
# subscription-manager repos \
    --enable=rhel-6-server-rh-common-rpms
```

For Red Hat Enterprise Linux 7, run:

```
# subscription-manager repos \
    --enable=rhel-7-server-rh-common-rpms
```

1.3.1. Content Delivery Network Channels

You can install Red Hat Enterprise Linux OpenStack Platform 5 through the Content Delivery Network (CDN). To do so, configure **subscription-manager** to use the correct channels.

Run the following command to enable a CDN channel:

```
# subscription-manager repos --enable=[reponame]
```

Run the following command to disable a CDN channel:

```
# subscription-manager repos --disable=[reponame]
```

Red Hat Enterprise Linux 6

The following tables outline the channels for Red Hat Enterprise Linux 6.

Table 1.1. Required Channels

Channel	Repository Name
Red Hat Enterprise Linux 6 Server (RPMS)	`rhel-6-server-rpms`
Red Hat OpenStack 5.0 (RPMS) for Server 6	`rhel-6-server-openstack-5.0-rpms`

Table 1.2. Optional Channels

Channel	Repository Name
RHEL Server Load Balancer (v6 for 64-bit x86_64)	`rhel-lb-for-rhel-6-server-rpms`
Red Hat Enterprise Linux 6 Server - Optional	`rhel-6-server-optional-rpms`

Red Hat Enterprise Linux 7

The following tables outline the channels for Red Hat Enterprise Linux 7.

Table 1.3. Required Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server (RPMS)	`rhel-7-server-rpms`
Red Hat OpenStack 5.0 for Server 7 (RPMS)	`rhel-7-server-openstack-5.0-rpms`

Table 1.4. Optional Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server - Optional	`rhel-7-server-optional-rpms`

Red Hat Enterprise Linux OpenStack Platform Installer

The following tables outline the channels for the Red Hat Enterprise Linux OpenStack Platform installer.

Table 1.5. Required Channels

Channel	Repository Name
Foreman-RHEL-6	`rhel-6-server-openstack-foreman-rpms`
Red Hat Enterprise Linux 6 Server (RPMS)	`rhel-6-server-rpms`
Red Hat Software Collections RPMs for Red Hat Enterprise Linux 6 Server	`rhel-server-rhscl-6-rpms`

Disable Channels

The following table outlines the channels you must disable to ensure Red Hat Enterprise Linux OpenStack Platform 5 functions correctly.

Table 1.6. Disable Channels

Channel	Repository Name
Red Hat CloudForms Management Engine	`"cf-me-*"`
Red Hat CloudForms Tools for RHEL 6	`"rhel-6-server-cf-*"`

Channel	Repository Name
Red Hat Enterprise Virtualization	`"rhel-6-server-rhev*"`
Red Hat Enterprise Linux 6 Server - Extended Update Support	`"*-eus-rpms"`

1.3.2. Red Hat Network (RHN) Channels

You can install Red Hat Enterprise Linux OpenStack Platform 5 through Red Hat Network (RHN).

Run the following to add a channel via RHN:

```
# rhn-channel --add --channel=[reponame]
```

Run the following to remove a channel via RHN:

```
# rhn-channel --remove --channel=[reponame]
```

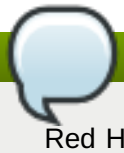

Note

Red Hat Network is only available via Red Hat Satellite on Red Hat Enterprise Linux 7.1. For more information on this product, see:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Satellite/6.0/html/User_Guide/index.html

Red Hat Enterprise Linux 6

The following tables outline the channels for Red Hat Enterprise Linux 6.

Table 1.7. Required Channels

Channel	Repository Name
Red Hat OpenStack 5.0 for RHEL 6 Server x86_64	`rhel-x86_64-server-6-ost-5`
Red Hat Enterprise Linux Server (v6 for 64-bit AMD64 / Intel64)	`rhel-x86_64-server-6`

Table 1.8. Optional Channels

Channel	Repository Name
RHEL Server Load Balancer (v6 for 64-bit x86_64)	`rhel-x86_64-server-lb-6`
RHEL Server Optional (v. 6 64-bit x86_64)	`rhel-x86_64-server-optional-6`
MRG Messaging v2 (for RHEL 6 Server x86_64)	`rhel-x86_64-server-6-mrg-messaging-2`

Red Hat Enterprise Linux 7

The following tables outline the channels for Red Hat Enterprise Linux 7.

Table 1.9. Required Channels

Channel	Repository Name
Red Hat OpenStack 5.0 for Server 7	**rhel-x86_64-server-7-ost-5**
Red Hat Enterprise Linux Server (v7 for 64-bit AMD64 / Intel64)	**rhel-x86_64-server-7**

Table 1.10. Optional Channels

Channel	Repository Name
Red Hat Enterprise MRG Messaging v. 2 (for RHEL 7 Server x86_64)	**rhel-x86_64-server-7-mrg-messaging-2**

Red Hat Enterprise Linux OpenStack Platform Installer

The following table outlines the channels for the Red Hat Enterprise Linux OpenStack Platform installer.

Table 1.11. Required Channels

Channel	Repository Name
Red Hat OpenStack 5.0 for RHEL 6 Server x86_64	**rhel-x86_64-server-6-ost-5**
RHEL OpenStack Platform Installer for Server 6	**rhel-x86_64-server-6-ost-foreman**
Red Hat Software Collections 1 (RHEL 6 Server x86_64)	**rhel-x86_64-server-6-rhscl-1**

1.4. Product Support

Available resources include:

Customer Portal

The Red Hat Customer Portal offers a wide range of resources to help guide you through planning, deploying, and maintaining your OpenStack deployment. Facilities available via the Customer Portal include:

- Knowledge base articles and solutions.
- Reference architectures.
- Technical briefs.
- Product documentation.
- Support case management.

Access the Customer Portal at https://access.redhat.com/.

Mailing Lists

Red Hat provides these public mailing lists that are relevant to OpenStack users:

- The **rhsa-announce** mailing list provides notification of the release of security fixes for all Red Hat products, including Red Hat Enterprise Linux OpenStack Platform.

  Subscribe at https://www.redhat.com/mailman/listinfo/rhsa-announce.

- The **`rhos-list`** mailing list provides a forum for discussions about installing, running, and using OpenStack on Red Hat based distributions.

  Subscribe at https://www.redhat.com/mailman/listinfo/rhos-list.

Chapter 2. Release Information

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality to be taken into consideration when deploying this release of Red Hat OpenStack.

Notes for updates released during the support lifecycle of this Red Hat OpenStack release will appear in the advisory text associated with each update or the *Red Hat Enterprise Linux OpenStack Platform Technical Notes*. This document is available from the following page:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform

2.1. Enhancements

BZ#895042

This update adds the ability to set the password for the keystone_admin user when running PackStack in interactive mode.

BZ#895586

Previously, when the status of an instance updated to 'Error', there was no feedback displayed on the UI. With this update, information about the error displayed for 'nova show [instance_id]' is displayed on the Instance Details UI page.

BZ#914648

This update introduces support for generating and distributing SSH keys to provide support for migrating instances via connections using QEMU and SSH to remote hypervisors. This update also ensures QEMU migration ports are open for Nova compute nodes and that Nova passes 'no_verify' to libvirt, making it possible for Nova to access compute nodes correctly when resizing instances.

BZ#970098

Previously, the service catalog used to return all endpoints, regardless of status. This meant that disabled endpoints were displayed as well.

Now, only enabled endpoints are returned by default.

BZ#978507

Compute users can now define a server group with an associated policy; supported policies are affinity and anti-affinity. When servers are created, they can be associated with a server group. When planning the new server, Compute enforces the affinity or anti-affinity policy among all instances associated with this group.

BZ#982689

A new feature has been added to Block Storage which allows users to delete multiple snapshots at once.

Users can now issue snapshot-delete for multiple snapshot ids at once.

The format of the command is: cinder snapshot-delete [uuid] <uuid...>

BZ#984270

With this release, you can now modify a given volume's type. When modifying a volume's type, the Block Storage scheduler checks if the volume's current host can accept the new type: as in, the scheduler checks if the host passes the filters when using the specified type. If the current host is suitable, the volume's corresponding manager then calls the right driver to change the volume's type accordingly.

If the current host cannot accept the new type, or if the volume driver is unable to modify the volume's type, then you will need to migrate the volume in order to change the volume type. Specifically, you will need to create a new volume of the type you want, and then migrate the contents of original volume to this one.

BZ#985500

With this release, you can now set a volume to read-only access. This feature allows you to give multiple users shared, secure access to the same data.

BZ#986680

Wildcards are now supported in Block Storage (cinder) for resource searches. For example, users can now use wildcards in the 'cinder list --display-name' or the 'cinder snapshot-list --display-name' commands.

BZ#986977

With this update, a case-insensitive search for resources is now possible in Block Storage (cinder). For example, the user can now run 'cinder list --status' with 'Available' instead of 'available'.

BZ#1006736

Previously, live-migration for the Compute service was only possible using the command line interface. There was no GUI option. With this enhancement, a new 'Live migrate' option is displayed in the menu for each instance, on the Instances page when logged in as an admin.

BZ#1017210

This update provides support for installing MariaDB on Red Hat Enterprise Linux 7 using the puppet-mysql module.

BZ#1026202

With this enhancement, Block Storage backup support has been added to the NFS driver.

BZ#1031689

Previously, as per most other OpenStack projects, Glance used to rely on an RPC library coming from oslo-incubator. This library graduated from the incubator and was moved into its own project, oslo.messaging. This request for enhancement tracks Glance's adoption of oslo.messaging.

Although the messaging library maintains backwards compatibility with older configuration options, upgrading the configuration files to use the `transport_url` and the new RPC options as necessary is highly recommended. The older configuration options are likely to be removed in future versions of the project.

BZ#1033117

A new enhancement to create and edit Availability Zones using the Dashboard has been added with this release. This is a useful Compute service feature to manage groups of Compute nodes. Users can now manage Availability Zones, a property of the Host Aggregates, using the new Host Aggregates panel in the Admin Dashboard.

BZ#1033132

A new enhancement to create and edit Host Aggregates using the Dashboard has been added with this release. This is a useful Compute service feature to manage groups of Compute nodes. Users can now create, edit and delete host aggregates using the new Host Aggregates panel in the Admin Dashboard, as well as add and remove hosts from an aggregate.

BZ#1035790

If the Block Storage service endpoint was set to Cinder v2, Dashboard displayed a 500 error when trying to display volume-related pages. With this update, Cinder v2 is supported and the OPENSTACK_API_VERSIONS dictionary in the local_settings file can now take a 'volume' attribute that can be set to either 1 or 2, depending on the Block Storage version the administrator wants to use.

BZ#1035896

You can now reset the state of multiple volumes in one call. To do so, specify the volume names in 'cinder reset-state'; for example, to reset volumes named 'vol1', 'vol2', and 'vol3' at the same time, run:

cinder reset-state vol1 vol2 vol3

This avoids having to reset each state individually.

BZ#1038668

Watchdog support has been added to the Libvirt driver. The watchdog device used is "i6300esb", and is enabled by setting the "hw_watchdog_action" property in the image properties or flavor extra specifications ("extra_specs") to a value other than "disabled".

Supported "hw_watchdog_action" property values, which specify the action for the watchdog device to take in the event of an instance failure, are "poweroff", "reset", "pause", and "none".

BZ#1040599

With this update, a Compute deployment which has a configured database slave (slave_connection) can send reads from periodic tasks to this slave. Periodic tasks typically result in consistently high loads to deployments. To improve performance, these tasks can now be offloaded to database slaves.

BZ#1040993

The Compute service determines what action to take when instances are found to be running that were previously marked deleted based on the value of the "running_deleted_instance_action" configuration key. A new "shutdown" value has been added to the list of configurable actions.

Using this new value allows administrators to optionally keep instances found in this state for diagnostics while still releasing the run-time resources.

BZ#1041014

Compute services are now able to shutdown gracefully by disabling the processing of new requests when a service shutdown is requested but allowing requests already in process to complete before terminating.

BZ#1041017

Notifications are now generated when a Compute host is enabled, disabled, powered on, shut down, rebooted, put into maintenance mode, or taken out of maintenance mode.

BZ#1041018

The Compute API now exposes the hypervisor IP address, allowing it to be retrieved by administrators using the "nova hypervisor-show" command.

BZ#1041023

Notifications are now generated upon the creation and deletion of keypairs.

BZ#1041026

The Libvirt driver now allows instance configuration to use video drivers other than the default (cirros), so that different video driver models and amounts of video RAM can now be specified. These values are configured by setting the "hw_video_model" and "hw_video_ram" properties in the image metadata. Currently supported video-driver models are "vga", "cirrus", "vmvga", "xen", and "qxl".

BZ#1041031

With this update: * Weights have been normalized in OpenStack Compute so that there is no need to artificially inflate multipliers. The maximum weight that a weigher puts for a node is 1.0, and the minimum is 0.0. * A new multiplier option, 'offset_weight_multiplier' (nova.cells.weights.weight_offset.WeightOffsetWeigher), has been introduced. * Stacking flags for weighers have been introduced. Negative multipliers should not be using for stacking, but the weighers are still compatible (they issue a deprecation warning message).

BZ#1041038

The /etc/nova/nova.conf configuration file has been updated to ensure that all configuration groups in the file use descriptive names. A number of driver-specific flags, including those for the Libvirt driver, have also been moved to their own option groups.

BZ#1041051

The Compute service now uses the tenant identifier instead of the tenant name when authenticating with OpenStack Networking (neutron). This improves support for the OpenStack Identity API v3, which allows non-unique tenant names.

BZ#1041053

With this update, the API call for attaching volumes to instances in Compute (servers/[INSTANCE_UUID]/os-volume_attachments) now accepts two additional parameters in the body: disk_bus and device_type. If these parameters are specified, the libvirt driver attempts to honor them when attaching the volume. The following values are accepted: * disk_bus: 'scsi' and 'virtio' * device_type: 'disk', 'cdrom', 'floppy', and 'lun'

BZ#1041055

The V3 API admin_actions plugin has now been separated into logically separate plugins so that operators can enable subsets of the functionality currently present in the plugin.

BZ#1041067

VMware Compute drivers now support a virtual-machine diagnostics call. Diagnostics can be retrieved using the "nova diagnostics INSTANCE" command, where INSTANCE is replaced by an instance name or instance identifier.

BZ#1041084

Transient database-connection failures are now recovered automatically. There are a variety of circumstances which can cause a transient failure in database connection (for example, the restart or upgrade of the database, migration of VIP between an HA pair, or a network failure). Compute now catches these "db-has-gone-away" errors by automatically reconnecting and retrying the last operation in such a way that the caller is able to continue whatever operation was in progress. The user no longer has to abort long-running operations (such as 'nova boot' or 'glance image-create') just because of a momentary interruption in database connectivity.

BZ#1041103

The Compute API now exposes a mechanism for permanently removing decommissioned compute nodes. Previously, decommissioned nodes would continue to be listed even if the compute service had been disabled and the system re-provisioned. The removal functionality is provided by the "ExtendedServicesDelete" API extension.

BZ#1041118

A new scheduler filter, "AggregateImagePropertiesIsolation", has been introduced. The new filter schedules instances to hosts based on matching namespaced image properties with host aggregate properties. Hosts that do not belong to any host aggregate remain valid scheduling targets for instances based on all images.

The new Compute service configuration keys "aggregate_image_properties_isolation_namespace" and "aggregate_image_properties_isolation_separator" are used to determine which image properties are examined by the filter.

BZ#1041487

The Block Storage service now sends notifications for attach and detach events, allowing other openstack services (e.g. Telemetry) to listen for and display the results to the user.

With these notifications, a volume's status can now be updated automatically in the Telemetry service. This, in turn, allows an administrator to search samples for volume status history.

BZ#1041489

With this release, you can now export the metadata of a volume backup. This ability provides a more complete way to backup and restore a Cinder volume. Information like Glance metadata can now be included in volume backups.

BZ#1041491

Previously, the 'cinder absolute-limits' command only displayed the maximum usable limits of a tenant. With this release, the same command now returns a tenant's consumed resources as well, namely:

* totalVolumesUsed (as in, the total number of volumes used) * totalGigabytesUsed

BZ#1041499

This update adds a new option for the Block Storage API service, namely osapi_volume_workers. This option allows you to specify the number of API service workers (or OS processes) to launch for openstack-cinder-api.

Typically, this option can be used to set the number of OS processes to the number of CPU cores/threads on a machine; doing so will greatly increase the number of API requests that can be handled per second. To set this option, run the following command on the openstack-cinder-api host:

openstack-config --set /etc/cinder/cinder.conf \ DEFAULT osapi_volume_workers [X]

Replace [X] with the target number of OS processes you wish to set. By default, the Block Storage API service will still run in one process (that is, osapi_volume_workers is set to None).

BZ#1041511

This release improves deployment adaptability by now allowing you to change the hostname used in the Domain Remapping feature (via cname_lookup). In doing so, CDN links can persist through changes in either hosting structure or providers.

With this, you can now freely use the CNAME feature of DNS in the presence of domain_remap middleware.

BZ#1041638

The OpenStack Block Storage service now features a Fibre Channel Zone Manager. This allows OpenStack Block Storage to automatically manage fibre channel SAN zoning, making it easier to deploy a properly-configured Block Storage fibre channel setup.

For more information on how to use the Fibre Channel Zone Manager, refer to the Red Hat Enterprise Linux OpenStack 5 Configuration Reference Guide.

BZ#1041662

With this release, you can now export and import volume metadata. This, in conjunction with the ability to backup and restore volumes, now allows you to restore a volume even in the event of a catastrophic database failure.

In addition, volume backup metadata support also adds portability to volume backups. Now, exporting a volume backup's metadata allows you to restore the volume backup on a different Block Storage service (or even a different cloud service altogether).

For more information, refer to:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/5/html/Cloud_Administrator_Guide/sect volumes.html#volume-backup-restore-export-import

BZ#1041693

A trial run of taskflow is implemented in cinder for only the create-volume task. The user should see no difference in behaviour.

With this implementation, cinder developers can now evaluate the usefulness and maintainability of taskflow actions without migrating the entire codebase.

BZ#1041696

The Chance and Simple schedulers are now deprecated. To replicate their behaviour, use the FilterScheduler instead. Both behaviours use the following setting in common:

the FilterScheduler instead. Both behaviours use the following setting in common:

scheduler_driver = cinder.scheduler.filter_scheduler.FilterScheduler
scheduler_default_filters = [`AvailabilityZoneFilter', `CapacityFilter',

`CapabilitiesFilter']

To replicate the behaviour of the Chance scheduler, add the following parameter:

scheduler_default_weighers = `ChanceWeigher'

To replicate the behavior of the Simple scheduler, use the following instead:

scheduler_default_weighers = `AllocatedCapacityWeigher'
allocated_capacity_weight_multiplier = -1.0

BZ#1041709

The Block Storage LVM driver now supports the LIO iSCSI target as an iSCSI back-end for an OpenStack implementation on Red Hat Enterprise Linux 7. This is required due to the move from scsi-target-utils (tgtd) to LIO in Red Hat Enterprise Linux 7.

BZ#1041737

Previously, it was not possible to mount all disks to a single directory when configuring multiple NFS servers as a backend using the filesystem store. This was due to the filesystem store only allowing administrators to configure a single directory using the filesystem_store_datadir parameter in the glance-api.conf file.

While it is possible to use MHDDFS (a FUSE plug-in: https://romanrm.net/mhddfs), which mounts multiple NFS servers to a single directory, MHDDFS does not allow you to evenly store the data on all the disks. Another major drawback is that it is very difficult to know the number and type of images stored on a disk when one of the disks is broken because the Glance registry stores the location specified in the filesystem_store_datadir parameter.

This enhancement fixes the above issues by adding multi-filesystem support to the current filesystem store.

BZ#1041738

Multi-location support allows Glance to store locations pointing to copies of the same image data stored in different places. Previously, this feature did not allow Glance to choose the store from which to download an image. This update adds two strategies for selecting the location to send back to the client - the location from which to download the image. The two strategies are:

1. Location order: A simple round-robin that goes through all the available locations in order - as they were inserted.

2. Store type: Allows users to specify store preferences. A possible combination is: `http, file`. This tells Glance to try to download images from HTTP stores before downloading them from file stores.

BZ#1041747

This update incorporates common code for working with databases such as session management, connections, engines, models, migrations and other utilities.

BZ#1041820

This update adds support for making new requests to Swift to download the remainder of

an image when downloading an image from Swift fails before the full image is transferred. The download operation is attempted a number of times based on the value set in the new configuration option 'swift_store_retry_get_count'. If the value of this key is set to '0', no attempts to retry downloading the remainder of the image are attempted.

BZ#1041859

Previously, users could not update their own passwords using the V3 API, only administrators could update users' passwords using the V3 API.

Now that the V3 API is the default, (no longer the V2 API), users can update their own passwords too.

BZ#1041860

A region resource has been introduced to the Identity API for constructing a hierarchical container of groups of service endpoints.

Previously, service endpoints could refer to a region as an arbitrary string. Regions can now be explicitly defined and managed through the Identity API, which are then referred to when adding service endpoints. This allows more control over region management.

BZ#1041863

When requesting validation of an Identity token, the ability to opt-out of including the service catalog in the response is now available to the requestor.

Identity tokens sizes can get large due to the inclusion of the service catalog. When the service catalog is not needed, it may be desirable to get a token that omits the service catalog. This ability was previously available only when generating a new token, but is now also available when validating of existing tokens.

BZ#1041864

Previously, Identity only provided API error and exception messages in English. Providing translated messages would be more useful when the requestor is using a different locale.

Now, translated messages are provided if a translation exists for the requestor's locale.

BZ#1041865

A change was introduced to limit the number of results returned by list commands.

This was done to avoid problems when the number of results can be larger than available memory. For example, a user list on a system where there are millions of users would return all of them.

Now users can provide a configuration value (list_limit=[integer value]) on calls to the Identity API list functions. Lists that have values beyond those limits will return a 'truncated' value in the body of the results.

BZ#1041875

Audits are essential for maintaining security of a system, and especially applicable to Identity since Identity is a security-focused service.

OpenStack has adopted CADF as the format for audit events. The Identity service now emits CADF events upon Identity and Token operations.

BZ#1041930

Implementers of custom Identity extensions may need to perform tasks when Identity resources such as users or groups are created, updated, or deleted. The ability to register callbacks for these events has been added to allow for more complex custom Identity extensions.

BZ#1041959

Previously, Identity trusts allowed the trust to be used to issue tokens for an unlimited number of times as long as the trust was valid.

This new feature adds the ability to specify the exact number of times that a trust can be used to issue tokens, allowing for uses such as a one-time use trust.

BZ#1041965

With this update, Role Based Access Control (RBAC) support has been added for Block Storage service. After copying the Block Storage policy.json file to /etc/openstack_dashboard/cinder_policy.json, some action such as Delete, etc. are displayed only when allowed by the policy.

BZ#1041976

With this update, a new look and feel was introduced for the dashboard featuring a horizontal instead of vertical navigation bar. This significantly differs from the new accordion navigation bar used upstream.

BZ#1041977

To add features or new components to the Dashboard, changes to the local_settings file were necessary. With this update, a plugin architecture for horizon is introduced. As a result, users can now place a config file in the /openstack_dashboard/enabled and restart the web server. For more information, see http://docs.openstack.org/developer/horizon/topics/settings.html#pluggable-settings-for-dashboards.

BZ#1041981

With this update, support for public container is added. This allows the user to set ACL as 'read' for the container to either public or private. This can be set at the container creation time or can also be updated at a later time by the user.

BZ#1042060

Pseudo-folders are similar to folders in your desktop operating system. They are virtual collections defined by a common prefix on the object's name.

With this update, users can create pseudo directories using the Create Pseudo-folder button without uploading an object.

BZ#1042121

This release also features upstream improvements to the implementation of i18n Message. These changes apply codebase improvements that streamline translation-related tasks and Message object processing, among other things.

BZ#1042140

This enhancement adds HOT software config in Orchestration (heat). This is available as a resource for defining software configuration. As a result, software can be configured using a variety of third-party configuration management tools.

BZ#1042148

The Orchestration service now has a new resource for describing cloud-init settings. This resource allows you to customize cloud-init settings for a stack, rather than use the hard-coded cloud-init settings provided by Heat.

BZ#1042150

With this release, heat_keystoneclient now uses the Identity service's v3 API exclusively. This API version allows non-admin users to perform autoscaling and use wait conditions.

BZ#1042151

This new feature allows specification of load balancer pool members in stacks.

Previously all pool members had to be specified at load balancer creation time.

Now, pool members can be dynamically defined.

BZ#1042152

Management features have been added to count stacks, provide pagination, provide Orchestration build information, and request scoping policy.

These are fundamental features are required for administering Orchestration.

BZ#1042153

With this release, you can now run multiple Orchestration engines simultaneously with just one database back-end. This feature adds horizontal scalability to the Orchestration service.

BZ#1042154

With this release, the OS::Neutron::SecurityGroup resource is now supported. This resource allows you to reliably specify security groups for the OpenStack Networking service. These security groups provide IP security for instances.

BZ#1042155

This feature adds new Orchestration-native pseudo parameters, to provide Orchestration template authors with a native way to access stack name, ID and region when authoring Orchestration templates using the HOT syntax.

Previously there was no native way to access this data, but there were equivalent parameters available when using the cloudformation-compatible template syntax.

Now, Orchestration supports the following new Orchestration-native pseudo parameters: OS::stack_id, OS::stack_name and OS::region.

BZ#1042156

This enhancement provides update support for OpenStack Networking resources.

Previously stack-update was not consistently working in OpenStack Networking resources.

With this enhancement stack-update will work consistently in OpenStack Network resources.

BZ#1042157

A new feature has been added which allows a server to be rebuilt with the "nova-server-rebuild" operation.

This is required so that a server will be rebuilt with a new image if the image becomes updated.

Now, the server will have a new image available.

BZ#1042158

This enhancement supports better security policy control. It is required because providing a secure environment for the Orchestration runtime is essential.

Now, users can customize the security policies used by Orchestration.

BZ#1042161

This enhancement removes the requirement of an Authentication user for in-instance credentials. As a result, security is improved, and auto-scaling and waitconditions are available for non-administrative users. This behavioral change is expected to be transparent to the user.

BZ#1042162

With this update, support for parameter validation has been added. This enhancement prevents users from entering error in parameters. As a result, any parameters that do not pass the template parameter specification will fail to validate.

BZ#1042163

With this enhancement, users are able to abandon a stack in Orchestration and still retain the stack resources. As a result, stack resources remain after a stack abandon operation.

BZ#1042164

This enhancement implements an update policy for auto-scaling. This was added to provide a consistent auto-scaling experience to the user. As a result, an update policy is available to auto-scaling users.

BZ#1042167

The Orchestration service now allows you to configure the SSL settings of each client independently. This will allow you greater flexibility in addressing security requirements unique to each client.

BZ#1042171

This enhancement adds an Orchestration resource for managing the new Database as a Service (Trove) feature.

BZ#1042176

This enhancement allows the exclusion of certain resources from appearing in the default list. Cloud providers may not want certain resources to be loaded, for example, those in the AWS namespace. As a result, specified resources are not loaded by the cloud environment.

BZ#1042177

The Orchestration service now allows you to filter stack information. This will allow you to better review relevant stack information across many stacks.

BZ#1042179

This feature allows Orchestration build information to be queried via the REST API. As a result, administrators are able to determine the version of the Orchestration engine or API they are utilizing when accessing customer support. Consequently, the build information can be presented to the user. For security reasons, it is important to note that the build information must be manually set.

BZ#1042180

With this enhancement, get_attr get_param is able to retrieve extra attributes. This allows users to perform better data selection through the HOT specification.

BZ#1042186

This feature allows the creation of multiple identical resources from one resource specification. This capability is useful for manual scaling purposes. As a result, a scaling group allows the creation of a set number of virtual machines for a stack, for example.

BZ#1042187

A new feature has been added which implements OpenStack Networking metering resources.

This allows template authors to configure IP ranges and to assign a label to them. For example, setting two labels: one for the internal traffic, and the other one for the external traffic. Each label will measure the traffic for a specific IP range. Then, bandwidth measurement will be sent for each label to the Oslo notification system and could be collected by the Telemetry service.

This enhancement gives users the ability to define OpenStack Networking metering labels and rules in an Orchestration template.

BZ#1042188

A new feature has been added to implement floating IPs in Compute, using OS::Nova::FloatingIP.

Prior to this change, floating IPs could not be used effectively.

Now floating IPs can be used effectively.

BZ#1042191

This release adds a new Orchestration resource type called OS::Heat::RandomString. This resource type randomly generates a string that can be accessed via an attribute.

BZ#1042192

This feature allows user_data to be passed to Compute (nova) in cloud-init free form. This allows users to pass user_data free form rather than with alterations. As a result, user data is unaltered by Orchestration (heat).

BZ#1042193

This update allows a resource's support status to be specified. As a result, resources can indicate to users their support status as an indication of the level of maturity, and support they may expect.

BZ#1042196

This feature enables notifications for stack creation, and autoscaling events. Notification of these events is expected to be useful to management systems.

BZ#1042203

Previously, Orchestration returns the parameters from the template in an unorganized manner. This can result in important parameters being randomly pushed towards the bottom as well as logical groupings of parameters being random.

With this update, the UI presents parameters in a specific order in specific groups. As a result, the UI can parse the template and use the information for better presentation. This is achieved by adding of the top level sections and validation to ensure the content is valid HOT format.

BZ#1042207

This enhancement allows resources to be filtered by support and maturity status. This was added to enable users to view a resource plugin's support status. As a result, only a filtered subset of resources would be returned by the 'resource-types' API.

BZ#1042209

This feature allows Compute image parameter constraints to be specified.

Previously, there was no check to ensure that Compute image parameters matched a specified Compute flavor.

Now, only flavors that exist can be specified in Compute.

BZ#1042210

This enhancement adds logging for resource and stack IDs. As a result, diagnostic capabilities are improved.

BZ#1042214

This enhancement provides a mechanism to paginate stack-list output. A large number of stacks may make it difficult to achieve consistent output of the stack list. Consequently, REST API requests can use pagination for an improved user experience.

BZ#1042217

This enhancement returns the stack identifier once a stack has been created. This is expected to mitigate possible confusion of a stack's identifier, as Orchestration formats the identifier with a UUID. Consequently, the stack identifier is returned upon stack creation.

BZ#1042221

This enhancement adds 'stack-adopt', which allows Orchestration (heat) to manage existing resources as a complete stack. As a result, Orchestration is able to manage resources that were present prior to Orchestration's deployment.

BZ#1042245

This enhancement allows previews of stack creation. This allows users to estimate a stack's resource consumption. Consequently, a stack description is produced that displays the desired resources, but is not actually instantiated.

BZ#1042248

This enhancement adds image parameter contraints for Image Service. This was added due to the possibility of stack creation failing as a result of a specified image not being present in the constraint. Consequently, images that do not meet the constraint requirements are now rejected.

BZ#1042256

This enhancement adds support for address pairs in Networking ports. Prior to this update, only a single address could be added to a Networking port. As a result, multiple addresses can now be added to a Networking port.

BZ#1043717

With this update, Block Storage has the ability to 'extend' (that is, expand or resize) a volume using the Dashboard.

BZ#1045289

Support has been added for VMware instances which boot from an ISO image. Software licensing which prevents the distribution of modified software images is now supported. In particular, this enables compliant creation of Microsoft Windows instances. For more information, see: https://wiki.openstack.org/wiki/BootFromISO

BZ#1046761

This update adds a software configuration REST API. This allows a mechanism for operation of the software configuration feature. As a result, software configuration operates as intended.

BZ#1048174

Previously, the size of a Glance image would be incorrectly reported under certain circumstances. This was caused by differences in the file size and the virtual size of the image.

The size of an image can refer to either the size of the file or the size of the actual image, which may not be the same in cases such as qcow2. This update splits the current size attribute into two separate attributes: image_size and file_size. The former refers to the real size of the image and the later the actual size of the uploaded file. In most cases, both fields will hold the same value. However, there are also many cases in which this value will differ.

This update considers both attributes important for an image, but not required for that image to exist. The value of the image_size attribute allows users to know the actual size of the image and how much space is needed to use that image, such as Cinder block allocation. The value of the file_size attribute is necessary to support quotas, CLI progress bars, rate-limits, and metering, etc.

While the actual image size could be included as part of the image properties or meta data, it is important to remember that the image_size attribute is a first-class citizen attribute in most external tools. A discrete attribute will ease the consumption of its value from Nova, Cinder and other tools that rely on the value of that attribute.

BZ#1048670

Previously, the get_file function allowed retrieval of free-form file data into the Orchestration template at stack life cycle operations. A new get_file function is added and as a result, longer files can now be included without cut and paste.

BZ#1049724

A new feature to send notifications for autoscaling events has been introduced to raise the bar on notifications for Third Party management systems to integrate with Red Hat Enterprise Linux OpenStack Platform. As a result, all autoscaling events come in through notifications now.

BZ#1050203

This enhancement adds support for NFS backends in the 'Block Storage' host group. As a result, NFS backends for Block Storage (Cinder) are available in the 'Block Storage' host group.

BZ#1051998

A Volume retype for Storwize SVC driver has been added.

Volumes managed by the Storwize driver can be retyped now.

BZ#1052799

With this enhancement, the Libvirt compute driver now supports providing modified kernel arguments to booting compute instances from AMI images. Kernel arguments are retrieved from the "os_command_line" key in the image metadata (as stored in the Image service), if the key's value is provided; otherwise, the default kernel arguments are used.

BZ#1052807

The Identity service default token duration setting has been reduced to one hour.

The Identity service previously defaulted to a token duration of 24 hours, which would results in scalability problems due to a large number of tokens being persisted in Identity's token database for tokens that are most likely not in use any more.

Now, the number of tokens persisted in Identity's token database will be vastly reduced compared to the previous default setting, resulting in improved scalability.

BZ#1053719

Previously it was not possible to create all network types in OpenStack Networking.

This feature allows OpenStack Networking to create all variants of network types.

BZ#1053727

With this update, floating IP addresses now have an operational status: 'Active', 'Down', or 'Error'. The operational status can be viewed with Dashboard and the command-line interface.

BZ#1053734

With this enhancement, PackStack now consistently performs the installation of the sos, sos-plugins-openstack and rhos-collector packages on all hosts.

BZ#1055853

The Compute service now includes a caching scheduler driver to help improve scheduler performance. The caching scheduler uses existing facilities for applying scheduler filters and weights, and caches the resultant list of available hosts.

When a user request is passed to the caching scheduler, the driver attempts to schedule based on the list of cached hosts (and only uses non-cached information if the attempt fails).

BZ#1055855

The Libvirt Compute driver now supports adding a Virtio RNG device to compute instances to provide increased entropy. Virtio RNG is a paravirtual, random-number generation device, which allows the compute node to provide entropy to compute instances in order to fill their entropy pool.

The default entropy device used is /dev/random, however the use of a physical hardware RNG device attached to the host is also possible. The use of the Virtio RNG device is enabled using the hw_rng property in the metadata of the image used to build the instance.

BZ#1055856

Identity previously emitted notifications for create, update, and delete operations for user, group, role, and project resources.

This has been extended to so that Identity also emits notifications for create, update, and delete operations for trust resources.

BZ#1056388

With this update, when logged in as an admin, a new tab titled Host Aggregates is available. This can be used to create and manage host aggregates and availability zones.

BZ#1056875

A new feature has been added that allows the Identity service's log messages to be translated based on the system locale.

Previously Identity only provided log messages in English. Now, translated messages will be provided if a translation exists for the system locale.

BZ#1056878

With this enhancement, you can create a new volume as a copy of an existing volume using the Create Volume option on the Dashboard.

BZ#1057830

With this update, Role Based Access Control (RBAC) support has been added for Image service. It is now possible to configure access to images via the /etc/openstack_dashboard/glance_policy.json file.

BZ#1057831

With this update, the functionality to add and modify an object or its information was added to the Dashboard.

BZ#1058131

The python-swiftclient package is no longer used by Object Storage (swift), and hence does not need to be installed on storage nodes.

It is desirable for python-swiftclient to depend on other packages, but not desirable to add a number of transient dependencies to server nodes.

Now, the python-swiftclient: - can use HTTP client libraries such as Requests, - may be included into Python 3 applications freely, and - imports python-keystoneclient utilities.

BZ#1058444

The Libvirt driver now supports using VirtIO SCSI (virtio-scsi) instead of VirtIO Block (virtio-blk) to provide block-device access for instances. Virtio SCSI is a para-virtualized SCSI controller device designed as a future successor to VirtIO Block, and provides improved scalability and performance.

BZ#1058494

Previously, the owner of an image was a private property in Glance's image v2. This enhancement makes the owner a public property and sends it back to the client when information on an image is requested.

BZ#1058577

Previously, the S3 middleware was in the keystone repository and package. Hence, in order to use the S3 middleware, the user had to either copy a specific file out, or pull in the whole keystone package. Neither of these were good solutions.

Now, the S3 middleware is in the keystone-client. A user can now deploy the middleware separately from the server, using just the client code.

BZ#1059963

An enhancement has been made to use Oslo Messaging for notifications.

Previously, Identity used the incubated code, but messaging has graduated to a stand-alone library. Using the new library keeps Identity consistent with the rest of the OpenStack services in the way it handles notifications.

This ensures that Identity gets all bug fixes and feature enhancements.

BZ#1062022

Copy-on-Write cloning for RBD-backed disks is now supported. As such, the Compute service no longer needs to download a glance image to local disk and then copy it again to RBD. Rather, disks can now be efficiently created directly in the Ceph system, thereby speeding up the creation of instances and, in the process, saving bandwidth and disk space.

BZ#1062815

To aid administrative troubleshooting, Guru Meditation reports are now sent by the Compute Service upon receipt of the SIGUSR1 signal. The report is sent to stderr and has the following sections: * Package — Displays information about the package to which the process belongs, including version information. * Threads — Displays stack traces and thread IDs for each of the threads within the process. * Green Threads — Displays stack traces for each of the green threads within the process. * Configuration — Lists all configuration options currently accessible through the CONF object for the current process.

BZ#1063583

This enhancement allows Networking to use additional CPU cores for processing agent requests. This multiprocessing feature was added to mitigate CPU contention issues. Consequently, Networking is able to process more concurrent requests. To enable this feature, edit /etc/neutron/neutron.conf: Uncomment the 'api_workers' and 'rpc_workers' options to commit any number of CPU cores.

BZ#1063585

With this update, when logged in as an admin, Resource Usage tab in the System Panel has a new Daily Report tab. You can select a period and generate a report per Project. This is useful to know how much resources various projects are using across all services.

BZ#1064958

This update adds support for NFS backends on non-HA controller host groups. Consequently, NFS backends for Block Storage are available on host groups 'Controller (Neutron)' and 'Controller (Nova Network)'.

BZ#1065182

To help improve debugging, Block Storage service log messages now include OpenStack request IDs. These IDs provide an easy way to associate log items with requests from other services such as Compute.

BZ#1066035

Previously, as is with most OpenStack projects, Cinder used to rely on an RPC library coming from oslo-incubator. This library graduated from the incubator and moved into its own project, oslo.messaging. This RFE tracks Cinder's adoption of oslo.messaging.

Although the messaging library kept backwards compatibility with older configuration options, it is highly recommended to upgrade the configuration files to use the `transport_url` and the new rpc options, where needed. The older configuration options are likely to be removed in future versions of the project.

BZ#1067211

This enhancement adds callbacks that notify Compute when a VIF has been plugged in and is ready on the host. In addition, Compute is notified when a floating IP address has been assigned or removed from the VIF. This was added to prevent occurrences where instances started before the VIF was ready, which resulted in no IP address allocation received from the DHCP server. Compute receives the notifications from Networking by default. Compute will not boot the instance if the notification has not been received; this will also apply if the VIF was plugged, but the notification was not received. Networking now contains configuration values to notify Compute using the API that a VIF was plugged in, and that a floating IP address was assigned, updated, or removed. The default configuration is to notify Compute, however the settings for the Compute API are not set by default; operators need to ensure that the Compute parameters are correctly configured.

BZ#1067217

With this update, a new disk configuration option is added to the Dashboard. In the Launch an Instance window, there is a new Advanced Options tab for disk partitioning. You can select either Automatic or Manual option from the dropdown list.

BZ#1067230

This update provides Block Storage service administrators with the ability to remove any existing quotas set for a particular tenant.

BZ#1081230

A new mechanism has been introduced to facilitate reliable notifications from OpenStack Networking to Compute about port-status updates. An instance startup is now postponed until ports used to run it are ready to use.

Before this release, there was no way for OpenStack Networking to reliably notify Compute about port status updates. Because of this, Compute was unable to wait for ports to become

ready, and started instances before OpenStack Networking was ready. This often resulted in instance boot failures. Compute can now reliably start multiple instances without any boot failures when using OpenStack Networking.

BZ#1085281

With this release, Red Hat Satellite 5.6 and Red Hat Subscription Asset Manager 1.4 are now notified of guest instance status. This feature was added for subscription reporting purposes.

In previous releases, users had to configure and run virt-who on each compute node in order to report hypervisor status to Satellite 5.6 or Subscription Asset Manager 1.4. Katello_notification is a daemon that runs on the OpenStack controller node and notifies Satellite 5.6 or SAM 1.4 of guest status changes. This centralizes guest status notifications, so each hypervisor does not have to do its own reporting.

Installing the 'katello_notification' rpm will install the daemon. Configuration is done via the '/etc/katello/katello-notification.conf' file. The URL for the Satellite or SAM instance needs to be configured as well as the 'mgmt_server' option. Valid options are 'spacewalk' for Satellite 5.6 or 'katello' for Subscription Asset Manager 1.4. To start the daemon, run 'systemctl start katello_notification'.

Note that in most cases, the Compute service will need to be configured to emit messages that detail guest usage. For this, the following options need to be set in the '/etc/nova/nova.conf' configuration file:

'notification_driver=messaging'
'notification_topics=notifications,subscription_notifications'
'instance_usage_audit_period=hour' 'instance_usage_audit=true'

Guest instances and host-to-guest associations will appear in either Satellite 5.6 or Subscription Asset Manager 1.4. Logs are available in '/var/log/katello_notification.log'.

BZ#1091367

With this update, the Red Hat Access Plugin for Red Hat Enterprise Linux OpenStack Platform now uses the new modular Red Hat Access Angular UI framework.

BZ#1091536

This enhancement adds support for Dell EqualLogic as a backend for Block Storage (Cinder). As a result, Dell EqualLogic is available as a Block Storage backend on the controller host groups: 'HA All In One Controller', 'Controller (Neutron)', and 'Controller (Nova Network)'.

BZ#1094385

The Block Storage (Cinder) service now supports multiple concurrent back-ends. To enable this feature, set the following Host Group parameters to 'true':

* cinder_multiple_backends * cinder_backend_[backend type] (where [backend_type] is each back-end type you wish to enable)

You will also need to set values specific to the selected back-ends.

BZ#1096054

With this enhancement, openstack-status now supports new services like openstack-ceilometer-notification and neutron-metering-agent.

BZ#1098716

This update adds the ability for PackStack to install an L3 metering agent. Installation of the metering agent is controlled by a new command line option (--os-neutron-metering-agent-install) and a new Boolean in PackStack answer files (CONFIG_NEUTRON_METERING_AGENT_INSTALL), which install the metering agent on all nodes running L3 agents when enabled.

BZ#1103148

This PackStack enhancement adds Load Balance as a Service (LBaaS) when Openstack Networking (neutron) is selected. LBaaS is optional and not activated by default. Activate LBaaS using either the command line, or the answer file.

BZ#1104093

Modular Layer 2 (ML2) and VXLAN are now the default neutron networking driver/type for Red Hat Enterprise Linux OpenStack Platform 5.

ML2 is the new Networking core plug-in introduced in OpenStack's Havana release. Superseding the previous model of singular plug-ins, ML2's modular design enables the concurrent operation of mixed network technologies. Controllers should now be configured with ML2 as the OpenStack Networking (neutron) driver by default, and the default tenant network type is VXLAN.

The monolithic Open vSwitch and linuxbridge plug-ins have been deprecated and will be removed in a future release; their functionality has instead been re-implemented as ML2 mechanisms.

BZ#1104219

With this release, the Compute scheduler runs in Active/Active mode for high availability (HA) by default.

This is required because Active/Active mode in HA for the Compute scheduler allows better scaling than the Active/Passive mode.

If you wish to run the Compute scheduler in Active/Passive mode, simply set scheduler_host_subset_size=1.

BZ#1104709

The heat-keystone-setup-domain script has been ported to the heat-common package of Red Hat Enterprise Linux OpenStack Version 5 (Icehouse). This script is needed to provide domain creation support in the puppet-heat module.

BZ#1105218

This enhancement enables support for multiple instances of the Dell EqualLogic backend for Block Storage (Cinder). As a result, all Dell EqualLogic parameters are now arrays instead of single values, with the exception of 'cinder_backend_eqlx' which remains a true/false switch to enable/disable the EqualLogic backend. In addition, all EqualLogic parameter arrays are expected to have the same number of elements: the first elements of each array form properties of one instance of the EqualLogic backend, while the second elements of each array form properties of another instance. This pattern continues for subsequent elements.

BZ#1108155

With this enhancement, PackStack configures the Orchestration service (heat) to use trusts by default. For the Orchestration service to work with trusts, each user must have a role for delegation (by default, this role is 'heat_stack_owner').

BZ#1109250

This update ensures that the Audit daemon is installed and running on each node (Controller, Network, and Compute nodes).

BZ#1110443

Previously, Orchestration's revision build-info was returned with the default "unknown", for example "revision": "unknown", which was not very useful.

This has been fixed to reflect the RPM version information of the packages produced at build time. For example: "revision": "1.0"

BZ#1111705

Previously, PackStack required password authentication even when the root user was running an all-in-one installation. With this release, the root user is no longer required to authenticate for all-in-one installations.

BZ#1112770

With this update, a new package, erlang-sd_notify is included. This package is required for rebasing rabbitmq-server to version 3.3.5. As a result, rabbitmq-server now supports systemd nofity mechanism.

BZ#1114660

You can now configure Dell EqualLogic storage device settings through the Red Hat Enterprise Linux OpenStack Platform Installer.

BZ#1115100

Virtual serial console access has been added to OpenStack instances. A new package and service have been added to support this feature. The cloud administrator must install the openstack-nova-serialproxy package and start the openstack-nova-serialproxy service.

BZ#1117115

If OpenStack Networking (neutron) is enabled, Packstack now displays a warning if the Network Manager service is active on hosts.

BZ#1120333

Red Hat Enterprise Linux OpenStack Platform Installer provides a wizard-based installation for easily deploying Red Hat Enterprise Linux OpenStack Platform on a set of hardware. It builds upon Foreman's capabilities and makes production-grade installations easier.

Red Hat Enterprise Linux OpenStack Platform Installer is fully supported in this release. You can use this tool to deploy a provisioning environment; this provisioning environment can then be used to deploy Red Hat Enterprise Linux OpenStack Platform 5 on Red Hat Enterprise Linux 7.1.

BZ#1122687

Red Hat Enterprise Linux OpenStack Platform is now fully integrated with the Red Hat Enterprise Linux High Availability Add-On, to support highly available environments. This means that a cloud infrastructure can now be set up so that if one of its controller nodes

fails (for example, because its host goes down), the machine can be brought back up and no data is lost (protection against single point of failure).

With this release, the following High Availability modes are supported:

* Active/Active * Active/Passive

BZ#1134992

Previously, Nova would default to cold snapshots of instances. As a result, instances needed to be shutdown before the snapshot was taken.

With this enhancement, Nova now uses live snapshots by default. Instances remain powered on during snapshots, and the process is unnoticeable to the user.

BZ#1139413

With this enhancement, a new command 'service-delete' has been added to the nova client to allow disabling services through the nova CLI as opposed to manually editing the nova-services table.

BZ#1149425

The default policy used by the Identity service does not allow for administration tasks to be delegated down to domain administrators. To address this, the Identity service now features a new domain-aware policy example; this example can be used as a baseline policy when configuring domains in the Identity service.

BZ#1180608

RBD snapshots and cloning are now used for Ceph-based ephemeral disk snapshots. With this update, data is manipulated within the Ceph server, rather than transferred across nodes, resulting in better snapshotting performance for Ceph.

BZ#1201851

When providing an affinity of the anti-affinity server group for two or more VMs, Compute previously only checked when the instance was booted but not with instance migration. This meant that two VMs in the same anti-affinity group could end up on the same compute host. With this update, the groups policy of a migrating instance is now checked, so that the policy is not violated when migrating.

BZ#1082762

With this update, the Red Hat Access Plugin for Red Hat Enterprise Linux OpenStack Platform now uses the new modular Red Hat Access Angular UI framework.

BZ#1115624

This update adds support for making new requests to Object Storage (swift) to download the remainder of an image if downloading an image from Object Storage fails before the full image is transferred. The number of download attempts is specified in the new 'swift_store_retry_get_count' configuration option. If the value of this key is set to '0', no attempts are made to retry downloading the remainder of the image.

BZ#1130503

Previously, the OpenStack plugins were necessary to collect detailed information from an installed OpenStack platform instance.

With this update, the sos-plugin-openstack package includes OpenStack plugins for the

SOS data collection tool for Red Hat Enterprise Linux 6. As a result, users of sosreport on OpenStack can now obtain detailed support information using the additional plugins contained in the sos-plugin-openstack package.

BZ#1130589

With this update, if OpenStack Networking (neutron) is enabled, Packstack now displays a warning if the Network Manager service is active on hosts.

BZ#1131621

With this enhancement, the PackStack utility now installs the sos, sos-plugins-openstack, and rhos-collector packages on each host. These packages enable the log collector tool to include RHEL OpenStack Platform information immediately after deployment.

BZ#1132408

With this enhancement, PackStack will not only install Puppet and it's dependencies on all hosts, it will also update all necessary packages to the latest versions available in repos.

BZ#1145257

Virtual serial console access has been added to OpenStack instances. A new package and service have been added to support this feature. The cloud administrator must install the openstack-nova-serialproxy package and start the openstack-nova-serialproxy service.

BZ#1172054

This enhancement allows process monitor to observe externally-spawned processes in the L3 and DHCP agents. This feature was previously disabled by default. As a result, any failing dnsmasq or neutron-ns-metadata-proxy services are properly restarted within 60 seconds. This behaviour results in more resilient neutron-dhcp-agent and neutron-l3-agent services.

BZ#1183796

RBD snapshots and cloning are now used for Ceph-based ephemeral disk snapshots. With this update, data is manipulated within the Ceph server, rather than transferred across nodes, resulting in better snapshotting performance for Ceph.

BZ#1206644

With this enhancement, a new command 'service-delete' has been added to the nova client to allow disabling services through the nova CLI as opposed to manually editing the nova-services table.

BZ#978500

The host argument for the 'nova evacuate' command has been made optional. This means that the user no longer has to know the host destination, simplifying evacuation in the case of an unplanned failure.

BZ#1041054

Compute now automatically attempts a controlled shutdown for stop, rescue, and delete instance actions. If the controlled shutdown fails, Compute falls back to a forced shutdown.

BZ#1127405

When using nova-network with multiple networks, it is now possible to set the MTU, enable or disable DHCP, set the DHCP server, and indicate whether the network shares addresses

with other networks. Previously, it was not possible to set these parameters on a per-network basis, making it more difficult to use nova-network with multiple networks. With this update, administrators now have more flexibility with settings when using multiple networks with nova-network.

BZ#1160405

RBD snapshots and cloning are now used for Ceph-based ephemeral disk snapshots. With this update, data is manipulated within the Ceph server, rather than transferred across nodes, resulting in better snapshotting performance for Ceph.

2.2. Technology Preview

BZ#1097869

With this update, python-django-sahara package has been updated. As a result, the menu, submenus for Data Processing (sahara) are displayed in the Red Hat theme.

BZ#1115231

The Firewall-as-a-Service (FWaaS) plug-in adds perimeter firewall management to Networking. FWaaS uses iptables to apply firewall policy to all Networking routers within a project.

FWaaS is currently in technical preview; untested operation is not recommended.

BZ#1118045

OpenStack Sahara enables the fast provisioning and easy management of Hadoop clusters on OpenStack. Hadoop is used to store and analyze large amounts of data, which is usually unstructured but can be a combination of both complex and structured data.

OpenStack Sahara is provided with this release as a Technology Preview (on Red Hat Enterprise Linux 7.1). For information on how to install OpenStack Sahara, refer to:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/5/html/Installation_and_Configuration_OpenStack_Sahara_Installation.html

BZ#1127432

The openstack-sahara package is now rebased to stable release 2014.1.2. This rebase applies several updates, including:

* Fix to closing of HTTP session in Ambari plugin. * Fix to put_file_to_hdfs method in hdfs_helper. * Setting python hash seed to 0 in tox.ini. * Updated global requirements. * Fix to deleting transient cluster before job is started.

BZ#1152204

Trove (Database-as-a-Service) is now included in this release as a Technology Preview. This service allow susers to quickly and easily utilize the features of a relational or non-relational database without the burden of administrative overhead. With Trove, users and database administrators can provision and manage multiple database instances as needed.

For more information about Trove, refer to https://wiki.openstack.org/wiki/Trove.

BZ#1110347

OpenStack Compute includes the concept of Cells, provided by the nova-cells package, for dividing computing resources.

Cells are provided in Red Hat Enterprise Linux OpenStack Platform as a Technology Preview at this time. Fully supported methods for dividing compute resources in Red Hat Enterprise Linux OpenStack Platform include Regions, Availability Zones, and Host Aggregates.

2.3. Release Notes

BZ#894888

Support for SPICE remote console access was recently added to the Compute (Nova) and Dashboard (Horizon) services. The spice-html5 package required to support SPICE access is however not included in this release. As such SPICE remote console access remains unsupported at this time.

BZ#1024032

In OpenStack Compute, attaching volumes to instances is not supported for device names like '/dev/hd*' which will cause the bus to be defaulted to 'ide', which cannot be hot plugged.

Attaching volumes to running instances is only supported for virtio, so the device name needs to be similar to '/dev/vd*'

BZ#1083224

As of this release, RabbitMQ replaces QPid as the default message broker handling Remote Procedure Call (RPC) communications in Red Hat Enterprise Linux OpenStack Platform. QPid will still be supported.

BZ#1102404

With Red Hat Enterprise Linux OpenStack Platform 5, ML2 is the new default Networking plug-in. Open vSwitch will also be used as the default ML2 mechanism driver.

BZ#1125374

Controller nodes require at least 3GB of free disk space to function. The mongodb service may fail to start if less than 3GB of free disk space is available.

BZ#1134069

New PackStack parameters have been added for specifying a HTTP proxy for 'subscription-manager'. The new parameters are:

* CONFIG_RH_PROXY * CONFIG_RH_PROXY_PORT * CONFIG_RH_PROXY_USER * CONFIG_RH_PROXY_PW

Note that the parameters CONFIG_RH_USER and CONFIG_RH_PW have to be completed in order for PackStack to successfully set a HTTP proxy on hosts.

BZ#1136167

Red Hat Enterprise Linux OpenStack Platform is only supported for use with the libvirt driver (using KVM as the hypervisor on Compute nodes) or the VMware vCenter hypervisor driver.

Refer to https://access.redhat.com/knowledge/articles/744153 for more information regarding the configuration of the VMware vCenter driver.

Red Hat is unable to provide support for other Compute virtualization drivers including the deprecated VMware "direct-to-ESX" hypervisor, and non-KVM libvirt hypervisors.

BZ#1136168

Some packages in the Red Hat OpenStack software repositories conflict with packages provided by the Extra Packages for Enterprise Linux (EPEL) software repositories.

The use of Red Hat OpenStack on systems with the EPEL software repositories enabled is unsupported.

BZ#1158213

In High Availability environments deployed through the Red Hat Enterprise Linux OpenStack Platform Installer, SELinux no longer needs to be Permissive on the HA controller to allow VNC access. A bug that required this in previous releases has since been fixed.

BZ#1115625

The sysfsutils package is now a dependency of the Compute service. This package provides the Compute service with the ability to attach Block Storage volumes.

In previous releases, sysfsutils was not a hard dependency of the Compute service; this required users to manually install sysfsutils as well during manual Compute installations. With this release, installing Compute packages will automatically install sysfsutils as well.

BZ#1116877

In a non-HA deployment of OpenStack on Red Hat Enterprise Linux 6.5 that uses LVM as the backing store for Block Storage, the disk partitioning creates a 500MB boot and 100GB LVM set for root. The rest of the space is used for a physical volume that is used for a cinder-volumes volume group that will be use by Block Storage to create volumes.

BZ#1131866

With this upgrade, a new PackStack parameter called "CONFIG_UNSUPPORTED" has been introduced. Setting this parameter to "y" activates other parameters which are not supported in the current release. Currently there is only one unsupported parameter which will be activated, "CONFIG_STORAGE_HOST", which enables the installation of OpenStack Image (glance) and OpenStack Block Storage (cinder) to a separate host.

BZ#975014

In order for the Compute service's resize command to work when using the libvirt driver and attempting to resize between nodes (the default resize method), Compute users on the compute nodes must have permission to perform passwordless SSH to the other compute nodes.

To set this up, generate SSH keys for the Compute user on each compute node, and then add the generated keys from the other compute nodes to the ~/authorized_keys file for the Compute user on each compute node.

2.4. Known Issues

BZ#908355

The SQL backend for Identity records tokens. It does not have a timeout, and it does not automatically remove tokens once they are recorded.

As a consequence, the SQL database can run out of storage space.

As a workaround, Identity now includes a command to remove tokens, namely 'keystone-manage token_flush'. This process should be scheduled to run regularly via cron. It is recommended that this command be run approximately once per minute.

BZ#973333

If a configuration file contains invalid or otherwise unparseable data OpenStack services that attempt to read values from it will fail to start. While this is expected behaviour in this situation the service initialization scripts do not currently display any output indicating that there was a failure.

To work around this issue when changing configuration files restart the service as normal and then manually confirm that the service is running. For example:

service openstack-keystone restart Stopping keystone: [OK] Starting keystone: [OK] # service openstack-keystone status keystone (pid 12632) is running...

BZ#1056405

A new configuration option, 'api_workers', can be configured for the OpenStack Networking service. Setting api_workers to a value greater than 0 results in multiple processes handling API requests. However, because this feature is currently experimental, it is likely to result in unreliable behavior. It is recommended, therefore, that this option not be used at this time.

BZ#1069157

At present, policy rules for volume extension prevent you from taking snapshots of in-use GlusterFS volumes. To work around this, you will have to manually edit those policy rules.

To do so, open the Compute service's policy.json file and change "rule:admin_api" entries to "" for "compute_extension:os-assisted-volume-snapshots:create" and "compute_extension:os-assisted-volume-snapshots:delete". Afterwards, restart the Compute API service.

BZ#1082443

Changing the neutron.conf rpc_workers option's default value of '0' can cause the Networking service to fail.

BZ#1085124

In some rare and specific circumstances on a few systems, data corruption can occur when using NET_DMA.

In these cases, a call to dma_skb_copy_datagram_iovec() to perform the DMA does not deliver the data to the expected location.

As a workaround, the chip manufacturer recommends that NET_DMA be disabled on the latest upstream kernels. This can be done by blacklisting the ioatdma module. A kbase article is available at https://access.redhat.com/articles/879293.

The problem will not occur with NET_DMA disabled. Recent hardware optimizations have effectively obviated the advantages of using the ioatdma driver on modern platforms. Therefore, Red Hat also recommends disabling ioatdma on all platforms.

BZ#1112383

The domain and foreman URL values will not set the host's hostname. Before the Red Hat Enterprise OpenStack Platform Installer is started: * The hosts' hostname must be set * The /etc/hosts file should have an entry that matches the hostname of the host.

Additionally, the foreman URL should use a FQDN that can be resolved by the nodes to be provisioned.

BZ#1112634

When using NFS shared storage for Compute instance storage, Red Hat recommends that you mount the share with the noac or the lookupcache=none option to prevent NFS clients from caching file attributes (for details, see the NFS man page). This enables migration and resizing instances between compute hosts that use the shared storage, but with slight performance penalties.

In a future release of RHOS, this requirement may be removed; the Release Notes will be updated when it is safe to use NFS shared storage for the Compute instance store without enabling the noac or lookupcache=none option.

BZ#1115232

VPNaaS is not available in this release of Red Hat Enterprise Linux OpenStack Platform.

BZ#1133920

If an existing haproxy process was already running before installing and running LBaaS (Load-Balancing-as-a-Service), attempting to start LBaaS will fail. This typically happens when upgrading to Red Hat Enterprise Linux OpenStack Platform 5 with an existing LBaaS service.

To work around this, you will have to kill the running haproxy process and restart the LBaaS agent:

kill $(pgrep haproxy) # service neutron-lbaas-agent restart

BZ#1136171

When using Keystoneclient authtoken middleware with Swift, the option to use memcached directly doesn't work. For example, the following won't work:

[filter:authtoken] memcache_servers = 127.0.0.1:11211

Instead, for token caching within Swift, authtoken middleware must be configured to use Swift's cache:

cache=swift.cache

BZ#1147293

An upgrade of the Red Hat Enterprise Linux OpenStack Platform installer is not supported from 5.0/5.0.1 to 5.0.2. The installer must be newly installed to the new version.

BZ#1150784

In deployments using OpenStack Networking, concurrently booting more than 63 VMs will result in only 63 VMs booting; the remainder will fail with error: "VirtualInterfaceCreateException: Virtual Interface creation failed"

BZ#1141923

The Red Hat Enterprise Linux OpenStack Platform Installer does not support UEFI booting. As a result, when attempting to perform a PXE-boot on a UEFI host, the installer will fail to download the required PXE images.

To work around this, switch the host to 'legacy' mode for PXE booting.

BZ#1144034

During subnet creation of the Public API, Admin API, and Management networks, the Red Hat Enterprise Linux OpenStack Platform installer does not validate whether:

1) IPAM is set to 'DHCP' and boot mode 'dhcp' (provisioning network), or

2) IPAM is set to 'internal db' and boot mode set to 'static'

If IPAM is set to 'internal db', then the range of IP addresses must be greater than or equal to the number of controllers (3) plus the number of VIPs (10-12 depending on which network).

Because these settings are not validated, it is possible for subnets to be created without enough IPs to allocate enough VIPs. In such cases, generating VIPs could result in puppet errors.

To prevent this, you need to manually ensure that your API subnets (Admin, Public, Management) have enough available IP addresses to allocate enough VIPs.

2.5. Deprecated Functionality

BZ#1040985

In OpenStack Compute, the OS-DCF:diskConfig API attribute is no longer supported in V3 of the nova API.

BZ#1044262

In previous releases, the 'glusterfs_disk_util' option allowed users to choose between df or du to calculate space usage. This option is now deprecated.

BZ#1056381

File injection into VM images is now deprecated; instead, Red Hat recommends that you use the ConfigDrive and metadata server facilities to modify guests at launch.

File injection is now disabled by default in OpenStack Compute. To enable file injection, modify the inject_key and inject_partition configuration keys in /etc/nova/nova.conf and restart the Compute services. Note that the file-injection mechanism will probably be disabled in a future release.

BZ#1068691

Although the Compute API currently supports both XML and JSON formats, usage of the XML format is now deprecated. API support for XML will be retired in a future release.

BZ#1150002

The monolithic Linux Bridge and Open vSwitch plug-ins are being deprecated in this release, and will be removed in a future release. We recommend that you use the ML2 plug-in with the Open vSwitch driver instead.

Note that the official instructions for upgrading to Red Hat Enterprise Linux OpenStack Platform 5 also contains instructions on how to migrate to the ML2 plug-in:

https://access.redhat.com/articles/1168953 https://access.redhat.com/articles/1168993 https://access.redhat.com/articles/1169003

Chapter 3. Upgrading

For details on upgrading to Red Hat Enterprise Linux OpenStack Platform 5 from your previous version, refer to https://access.redhat.com/articles/1177953. This article also contains links to instructions for each recommended upgrade method.

Revision History

Revision 5.0.0-46 **Thu Apr 16 2015** **Don Domingo**
Updated for fourth maintenance release of Red Hat Enterprise Linux OpenStack Platform 5.

Revision 5.0.0-44 **Wed Mar 04 2015** **Andrew Dahms**
Corrected validation error.

Revision 5.0.0-43 **Fri Feb 27 2015** **Andrew Dahms**
Updated the list of supported browsers.

Revision 5.0.0-42 **Mon Feb 16 2015** **Andrew Dahms**
Added a known issue for the RHEL-OSP installer.

Revision 5.0.0-41 **Fri Feb 6 2015** **Don Domingo**
Fixed typos in subscription info.

Revision 5.0.0-40 **Wed Dec 3 2014** **Andrew Dahms**
Minor revision to fix a validation issue.

Revision 5.0.0-39 **Wed Dec 3 2014** **Andrew Dahms**
Updated the description of a known issue.

Revision 5.0.0-38 **Wed Dec 3 2014** **Don Domingo**
Updated for third maintenance release of Red Hat Enterprise Linux OpenStack Platform 5.

Revision 5.0.0-36 **Thu Nov 13 2014** **Andrew Dahms**
BZ#1151284 - Reformatted the description of software channels.

Revision 5.0.0-35 **Thu Nov 06 2014** **Don Domingo**
Updated for second maintenance release of Red Hat Enterprise Linux OpenStack Platform 5.

Revision 5.0.0-32 **Mon Oct 13 2014** **Don Domingo**
BZ#1150002: added reminder about deprecation of networking plug-ins.

Revision 5.0.0-30 **Tue Sep 30 2014** **Don Domingo**
Updated for first maintenance release of Red Hat Enterprise Linux OpenStack Platform 5.

Revision 5.0.0-29 **Wed Sep 24 2014** **Don Domingo**
Applied minor edits for readability.

Revision 5.0.0-28 **Wed Sep 3 2014** **Don Domingo**
Added links to upgrading documentation. This update coincides with several component rebases.

Revision 5.0.0-27 **Fri Aug 22 2014** **Don Domingo**
Added further updates to Red Hat Enterprise Linux OpenStack Platform 5.0 for both Red Hat Enterprise Linux 6.5 and 7.0. This includes full support for Red Hat Enterprise Linux OpenStack Platform Installer.

Revision 5.0.0-26 **Tue Aug 5 2014** **Don Domingo**

Added information relating to Red Hat Enterprise Linux OpenStack Platform Installer.

Revision 5.0.0-23 Thu Jul 24 2014 Don Domingo

Added further updates to Red Hat Enterprise Linux OpenStack Platform 5.0 for both Red Hat Enterprise Linux 6.5 and 7.0.

Revision 5.0.0-21 Fri Jul 11 2014 Don Domingo

BZ#1110347: added OpenStack Sahara to Technology Preview list.

Revision 5.0.0-20 Mon Jul 7 2014 Don Domingo

Red Hat Enterprise Linux OpenStack 5.0 (for Red Hat Enterprise Linux 7.0) GA release build.

Revision 5.0.0-16 Tue Jun 24 2014 Don Domingo

Applied new recommended publication settings based on new documentation portal features.

Revision 5.0.0-15 Mon Jun 23 2014 Don Domingo

Applied new recommended publication settings based on new supported build features.

Revision 5.0.0-14 Wed Jun 18 2014 Don Domingo

Edited title for better portal visibility.

Revision 5.0.0-13 Tue Jun 17 2014 Don Domingo

Applied recommended build settings as per updates to https://access.redhat.com/site/documentation/.

Revision 5.0.0-9 Tue May 13 2014 Don Domingo

Added FTP information and additional Known Issues for Beta release.

Revision 5.0.0-7 Mon May 12 2014 Don Domingo

Red Hat Enterprise Linux OpenStack 5.0 (for Red Hat Enterprise Linux 7.0) Beta build.

Revision 5.0.0-1 Fri May 02 2014 Don Domingo

Red Hat Enterprise Linux OpenStack 5.0 (for Red Hat Enterprise Linux 7.0) initial build.